UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE

AD-A261 487

(2)

REPORT DOCUMENTATION PAGE

DTIC ELECTE MAR 5 1993 S C D

1a. REPORT SECURITY CLASSIFICATION	
2a. SECURITY CLASSIFICATION AUTHORITY	3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; distribution unlimited.
2b. DECLASSIFICATION / DOWNGRADING SCHEDULE	
4. PERFORMING ORGANIZATION REPORT NUMBER(S)	5. MONITORING ORGANIZATION REPORT NUMBER(S) AFOSR-TR- 93 0089
6a. NAME OF PERFORMING ORGANIZATION Department of Mathematics University of Washington	6b. OFFICE SYMBOL (If applicable)
7a. NAME OF MONITORING ORGANIZATION AFOSR/NM	
6c. ADDRESS (City, State, and ZIP Code) Seattle, WA 98195	7b. ADDRESS (City, State, and ZIP Code) AFOSR/NM Bldg 410 Bolling AFB DC 20332-6448
8a. NAME OF FUNDING / SPONSORING ORGANIZATION AFOSR	8b. OFFICE SYMBOL (If applicable) NM
9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER AFOSR-89-0081	
8c. ADDRESS (City, State, and ZIP Code) AFOSR/NM Bldg 410 Bolling AFB DC 20332-6448	10. SOURCE OF FUNDING NUMBERS

PROGRAM ELEMENT NO.	PROJECT NO.	TASK NO.	WORK UNIT ACCESSION NO
61102F	2304		

11. TITLE (Include Security Classification)
Methods of Optimization Under Uncertainty: Final Scientific Report

12. PERSONAL AUTHOR(S)
R. T. Rockafellar

13a. TYPE OF REPORT	13b. TIME COVERED	14. DATE OF REPORT (Year, Month, Day)	15. PAGE COUNT
Final, Scientific	FROM ____ TO ____	1992/12/29	13

16. SUPPLEMENTARY NOTATION

17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
FIELD	GROUP	SUB-GROUP	large-scale optimization, stochastic programming, decision-making under uncertainty

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

Research under this grant has focused on large-scale optimization methodology connected with the solution of problems in which decisions must be made in the face of uncertainty: stochastic programming problems. The principal techniques developed for modeling such problems have been used various new kinds of decomposition into small-scale optimization problems in extended linear-quadratic programming. Extended linear-quadratic programming goes beyond ordinary linear and quadratic programming in allowing for objective functions to incorporate penalty terms and other features which create piecewise linear or quadratic formulas. The new decomposition techniques include primal-dual Lagrangian decomposition and forward-backward splitting.

In total, four-year grant supported the writing of 16 technical papers (12 already in print or aboutto be), the devlopment and documentation of 2 computer codes, and the completion of 3 doctoral dissertations.

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT	21. ABSTRACT SECURITY CLASSIFICATION	
☐ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT ☐ DTIC USERS		
22a. NAME OF RESPONSIBLE INDIVIDUAL Maj. James M. Crowley	22b. TELEPHONE (Include Area Code)	22c. OFFICE SYMBOL NM

DD FORM 1473, 84 MAR

83 APR edition may be used until exhausted.
All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE

UNCLASSIFIED

Final Technical Report on Grant AFOSR–89–0081

METHODS OF OPTIMIZATION UNDER UNCERTAINTY

November 1, 1988, to October 31, 1992

Prof. R. T. Rockafellar, Principal Investigator
Dept. of Mathematics GN–50
University of Washington
Seattle, WA 98195

December 29, 1992

SUMMARY OF RESEARCH AND PUBLICATIONS

Overview

Many important problems in management and engineering involve interactive decisions that must be taken in successive time periods and in the face of uncertainty. In logistics, for example, inventory systems have to be managed at adequate levels in a cost-minimizing manner despite vagaries in demand. Distribution systems have to be organized and programmed to deliver stocks to their destination in reasonable time even though random delays and breakdowns in transport are possible.

The uncertainty in these problems comes mainly from an inherent lack of full knowledge of what the future may bring. It can in some cases also reflect imperfect information on the present or past circumstances of the system being guided. Either way, there are formidable obstacles to optimizing so as to obtain the "best" decision policy for a given purpose. The difficulties are computational, because problems of enormous size can be generated in trying to take the possibilities of future branching adequately into account, but they are also conceptual. Practical ways of modeling the uncertainty, so as to get somewhere with it mathematically, have been much in need of development.

Until the last few years, there was little real hope of being able to optimize under uncertainty. For the most part, deterministic models were set up and utilized, even when stochastic elements were rampant. One notable exception was linear-quadratic regulator theory in stochastic control, which however covers a very particular situation in systems engineering, does not allow for any constraints, and has not proved amenable to generalization.

This lack of methodology has been unfortunate, because solutions to deterministic models of stochastic situations tend to be fragile. Decisions based on such models have no provision for hedging against eventualities that, although unlikely, could be serious if they arise. The consequences of neglecting uncertainty can therefore be worse than mere suboptimality, where less money is saved, say, than would be the case if the true solution were followed. They can be felt in a lack of built-in redundancy in the decision pattern, where too much can depend on quantities that, in the end, shouldn't be counted on.

A simple illustration in logistics would be a policy of depending entirely on one source of supply for a critical item, just because that source was slightly cheaper. In a deterministic world, nothing could be wrong with such a policy. But in the real world, something might happen to interrupt the source and cause a shortage of the item just when it is suddenly needed.

The best known optimization approach to dealing with uncertainty over time has for many years been that of *dynamic programming.* While initially attractive in theory, dynamic programming has proved unworkable in most applications with finite time horizon due to the "curse of dimensionality." Even if computers were able to use it to calculate solutions in problems of realistic size, however, there would be definite mathematical drawbacks to its use, giving motivation to look for something better.

First, dynamic programming suffers from a need to discretize in state space as well as in time and probability. This essentially means that many of the features of a problem that might potentially be valuable in solving it, like derivatives and convexity, are simply thrown away. Dynamic programming is also handicapped by its mode of working backward in time from the terminal period. This seems counter to the notion that the present should be more influential than the future, not only in influencing the nature of a solution but in finding a solution. It has the effect that if computations are cut off before they are finished, the output is useless.

Dynamic programming has furnished interesting "steady state" solutions to some problems over an *infinite* time horizon. But models in which such solutions make sense have a very special character, where randomness is associated with a known probability distribution that never changes over the entire future, and no goals are set up to be met other than a sort of stabilization of a given system. Such models are far removed from the problems under discussion here.

Quite a different approach to optimization under uncertainty has been building in the area of *stochastic programming.* The ambitions in dynamic programming of being able to encompass a vast spectrum of relationships between information, observation and the making of decisions, are waived in stochastic programming. The emphasis instead is on more specific structure supportive of solution techniques such as are inspired by the successes in linear-quadratic programming and convex programming.

The bulk of the work in stochastic programming has so far concentrated on the two-stage case. In this case, a decision that is to be made now, under constraints, will be followed by a single corrective decision after some aspect of the future becomes known. With some mathematical manipulation, the cost of the corrective decision, as a function of the initial decision z, takes the form of an expectation $c_2(z) = E\{\varphi(z, s)\}$, where s is the state of the future, a random variable. With the initial costs denoted by $c_1(z)$, the problem then comes down to minimizing $c_1(z) + c_2(z)$ subject to constraints on z.

Simple as this may look, the problem is numerically still very formidable due to the form of $c_2(z)$. If the expectation refers to integration with respect to a continuous probability distribution in several variables, it generally can only be approximated in some way, so that at best one obtains an approximation to $c_2(z)$ and $\nabla c_2(z)$ for any given z. There are various forms of approximation in which the continuous probability distribution is replaced by a "well chosen" discrete distribution concentrated in finitely many points so as to obtain upper or lower bounds. Other forms of approximation in two-stage stochastic programming rely on sampling of the probability distribution.

The kinds of problems of optimization under uncertainty that have stimulated the work under this grant are large-scale stochastic programming problems with dynamical structure generally extending over a number of future "stages." In all stochastic programming, the goal is to make a wise choice of a required *here-and-now* decision. Again, the difficulty is that this decision must be taken in advance of full knowledge of the realizations of certain random variables, such as the demands, system failures, or situational emergencies that might occur. Ordinary deterministic optimization assumes such knowledge and, by relying on this idea despite the realities, produces "fragile" decisions which could have unpleasant outcomes. Stochastic programming attempts to identify a more robust sort of decision by utilizing various representations of how the future might evolve, and then providing the mechanism that enables the here-and-now decision to hedge against negative eventualities but take advantage of positive ones.

Any representation of the future requires a high degree of simplification if a problem is to be kept manageable, but even a greatly reduced model can be far superior to a deterministic one. Until the last few years, most computational work in stochastic programming has in fact centered on two-stage models, where the here-and-now decision is supplemented by only by one subsequent opportunity for recourse. In contrast, the work under this grant has been aimed at pioneering the case where recourse actions will be possible at more than one future time, so that the problem is *multistage*. Necessarily then, the dynamical structure of decision-making becomes a key topic for analysis, even though, as always, the end product of the theory is just a well hedged *initial* decision.

To avoid taking on too many difficulties at once, the project has mainly been formulated in terms of problem models in the category of *extended linear-quadratic programming*. Mathematically, this refers to the use of linear constraints but objective functions that may be linear or quadratic, but could also just be *piecewise* linear or quadratic and thus able to incorporate standard types of penalty terms. (A failing of some past work in optimization under uncertainty was a treatment of all constraints as if they were black or white, instead of having gray shades which correspond to the invoking of penalty costs as desired values begin to slide. The concept and theory of extended linear-quadratic programming was developed by the P.I. under predecessor grants from AFOSR.)

On the computational level, therefore, it has been natural to look hard at large-scale problems of extended linear-quadratic programming in which a special dynamic and stochastic structure is present. A prime goal has been the discovery of features within such structure that can be used to decompose a large-scale problem iteratively into smaller tasks, and the numerical experimentation with algorithms based on such features. The efforts in this direction have focused on Lagrangian saddle point representations of optimality, which have revealed a number of new algorithmic possibilities.

As a natural counterpart, research has proceeded on how problems beyond the mold of extended linear-quadratic programming could be approximated sensibly by such problems in a local sense. This has involved the analysis of data perturbations and their effects on solutions. The perturbational results, utilizing nonsmooth analysis, have been applied in turn to questions of approximation that arise in replacing the true random variables in a problem by discrete variables generated through random sampling. This has led to a statistical theory of the behavior of optimal solutions in stochastic programming.

Taking part in the project, besides the P.I. himself, have been a number of the P.I.'s current or past Ph.D. students, as well as Roger Wets, a long-time collaborator. All told, the grant has supported the production of

- 12 technical articles now in print or soon to be
- 4 more research articles, one already submitted for publication, and three more as technical reports not yet in publication form
- 2 documented computer codes for new numerical methods of solution
- 3 doctoral dissertations completed, two others in the making. In addition, many new research results have been obtained that are still being augmented and will be written up in the near future.

Scenarios and Hedging

First on the list of publications to be described is "Scenarios and policy aggregation in optimization under uncertainty" [1], written with Roger Wets. This was put together under the predecessor AFOSR grant, but was substantially reworked and improved during the period being reported on here. The paper makes a very substantial contribution to the practical feasibility of techniques for optimization under uncertainty, and indeed, it has received much attention in the stochastic programming community.

The distinguishing feature in [1] is that a sophisticated statistical or probabilistic background for a given problem is not at all assumed. Rather, it is assumed only that the modeler can come up with a finite set of "scenarios" representing how the future may evolve and can describe how these scenarios branch, as well as supply guesses as to the branching probabilities. Where at present people simply solve the deterministic scenario subproblems corresponding to the different choices of the future, and then by nothing firmer than vague intuition try to come up an appropriate compromise not based so dangerously on optimizing from the perspective of a soothsayer, the paper shows how to iteratively modify such subproblems and aggregate their solutions so as to eventually create a policy that is optimal in a certain natural sense. It builds in this way on whatever solution technology is already available for the subproblems.

While the scenario hedging method in [1] is attractive from several angles, and is virtually the first algorithm designed directly for multistage, rather than merely two-stage, problems, its rate of convergence is slower than one would like. Therefore, efforts have been made to speed up convergence through supplementary devices. Paper [13], also written with Wets, has this aim. It makes improvements in terms of a kind of cutting plane approximation to the dual elements that are needed in representing the price of future information in the iterated subproblems.

Envelope Methods

Paper [2], "Computational schemes for large-scale problems in extended linear-quadratic programming," sets up a new framework for solving problems of finding a saddle point of a linear-quadratic convex-concave function on a product of polyhedral sets in spaces of high dimension. Finding such a saddle point is equivalent to solving an extended linear-quadratic programming problem along with its dual. The saddle point framework was shown in papers written by Roger Wets and the P.I. under the predecessor AFOSR grant to be a natural one for multistage stochastic optimization. Most of the literature on numerical techniques in this area has been aimed instead at purely primal or dual formulations reflecting the traditional paradigms of linear and quadratic programming with hard constraints, but because of this bias certain special features have been missed by others. It was not noticed that, in saddle point form, one is able to achieve an important simplicity of problem representation despite the use of penalty terms for constraints, which is a necessity often in the face of uncertainties. Furthermore, this simplicity can be gained in such a way that the Lagrangian function, for which one wishes a saddle point, is "doubly decomposable." This means that if one fixes either the primal or dual argument, the Lagrangian is highly separable in the other argument.

Paper [2] lays down the rules for exploiting this sort of double decomposability and, in the process, introduces a new class of so-called finite envelope methods. Such techniques are related to the finite generation methods in stochastic programming that were introduced earlier by Wets and the P.I. (again in research supported by AFOSR). but the latter required either the primal or the dual dimension to be low. This is appropriate only in the two-stage case of decision structure. however. For the new methods. convergence is established when certain line search steps are included. Line search appears feasible in consequence of the double decomposability.

The same themes are continued in the paper "Large-scale extended linear-quadratic programming and multistage optimization" [3]. The emphasis in this case is on the role of the dynamical structure and how to take advantage of it in ways other than the well trodden ones in mathematical programming, which involve sparsity patterns in large matrices.

In order to provide for numerical testing of finite envelope methods. a FORTRAN code was written by Stephen E. Wright. a Ph.D. student, and documented in [9]. (Wright is now at the T. J. Watson IBM Research Laboratories in Yorktown Heights. New York. where he is a key member of a team devoted to the development of stochastic programming.) The code was modularized so that parts could also be utilized and extracted for various other projects as well. It concentrated on problems with discretized dynamics. which made it possible readily to generate test examples with large numbers of variables. but nevertheless possessing inherent stability and solutions that readily could be verified.

Another Ph.D. student, Ciyou Zhu, developed the finite envelope idea further and showed it could lead to algorithms analogous to conjugate gradients or steepest descent, but able to cope with box constraints as well as the discontinuities of second derivatives that underlie the structure of extended linear-quadratic programming problems. Zhu made use of Wright's code [9] to test these algorithms numerically alongside of the basic finite envelope algorithms in [2]. He was able to solve large-scale dynamical problems with many time periods, involving as many as 100.000 primal and 100.000 dual variables. The test results have been presented in paper [10]. which also develops the theory behind the special algorithms.

These algorithms turned out to be superior to the basic ones. but all the finite-envelope algorithms were successful in tackling difficult problems whose structure has so far been rather neglected or poorly understood. Zhu's versions derive from a novel concept of projected gradient iterations pursued simultaneously in the primal and dual problems in such a way that massive decomposition can take place. A new form of "information feedback" between the primal and dual calculations leads to rather dramatic speedups. Even a version of the procedure that resembles steepest-descent, a first-order method, ends up behaving almost like a second-order method in its convergence properties. Something important seems to have been uncovered here, but the theoretical implications are yet to have been fully grasped in their potential for extension to other schemes.

For problems of extended linear-quadratic programming on a smaller scale, Zhu wrote a FORTRAN code independent of Wright's. This has been documented by Zhu and the P.I. in [11]. Zhu's dissertation [12] was completed in August, 1991. It lays down the theory behind his primal-dual projected gradient algorithms. It also includes a remarkable technique for accelerating the proximal point algorithm as an outer scheme to introduce strong convexity and stabilize the extended linear-quadratic subproblems.

Supporting work on the properties of extended linear-quadratic programming has been carried out in the P.I.'s papers [14] and [17]. The recent paper [15], "Lagrange multipliers and optimality," likewise falls in this category, but builds the foundations for approximating more general problems by ones of this type.

Perturbation and Approximation

The Ph. D. dissertation of Steve Wright [8], completed in December of 1990, grew out of his work with setting up code for our numerical experiments on decomposition using finite-envelope methods, as already described. It provides important theoretical support not only for this specific endeavor, but also for other algorithmic developments. Basically, the dissertation concerns the approximation of underlying infinite-dimensional problems (with continuous probability or continuous time) by discretized finite-dimensional problems, and the establishment of criteria under which the solutions to the discretized problems converge to one for the underlying problem as the approximations get finer.

This may sound like a traditional topic, but in the setting required here a major challenge is encountered. The core of the difficulty is that the approximations should not merely be in some abstract sense, but rather of a special form which our work had earlier identified as especially conducive to computations, namely one allowing for massive decomposition and parallelization. This means a *dual* as well as *primal* discretization, which goes beyond traditional thinking. Wright has been able in [8] to prove powerful theorems in this respect, and for such a purpose even had to do innovative studies on the frontiers of nonsmooth analysis in this area of optimization.

Another type of approximation has been pursued in the paper "Sensitivity analysis for nonsmooth generalized equations" [6], which was written by the P.I. with Alan J. King, a former Ph. D. student of his (1986) whose dissertation on the statistical properties of solutions to problems in stochastic programming was supported earlier by AFOSR. (King is now at the IBM Research Center, Yorktown Heights, and is charged with developing stochastic programming applications and software for IBM. He heads the group to which Steve Wright belongs, as mentioned above.) This paper concerns local approximation to the mapping that gives the optimal solution set in a problem as a function of the problem parameters. Such a mapping is unlikely to be differentiable, so it can't just be "linearized" for example. Instead, concepts of nonsmooth analysis must be used to discover the nature and properties of the kind of approximation that should be made.

The paper [6] with King has provided the theoretical underpinnings for an important advance in simulation techniques in stochastic programming. A formidable difficulty in computational approaches to stochastic programming is that of justifying the use of *sampling*. The random variables in a given problem of optimization may have complicated joint distributions, but typically they can at least be sampled empirically or through computer simulation. In that way, one gets a discrete empirical distribution, and this can be thought of as providing an approximation to the given problem. Since the results of sampling are themselves random, the approximate problem is in a sense random, and so then is its solution. The question then arises as to the statistical properties of this random solution.

For instance, as the sample size increases, can one count on the distribution of the random solution concentrating more and more around the true solution to the given problem? A particularly tantalizing goal would be to understand this question well enough to give guidelines in advance as to the size of the sample that should be taken, so as to be sure of a specified degree of statistical confidence in the result of solving the approximate problem. To get anywhere with this, a broader form of asymptotic statistical theory must be developed, and this requires the analysis of sensitivity to perturbations in the case of certain kinds of generalized equations that serve as the optimality conditions in stochastic optimization.

Article [7], also written with King and entitled "Asymptotic theory for generalized M-estimation and stochastic programming," goes a long way toward this goal. It tackles the main issue and obtains results on the generalized differentiability of the solution mappings with respect to parameters on which the equations depend. Central limit properties are obtained that fit the requirements, even though classical statistical theory isn't applicable.

Crucial as backup for this statistical work, in particular in establishing the special forms of approximate but nonnormal distributions that come up, has been the theoretical contribution of the P.I. in [5].

A paper with ideas to those in [6], but in a direct framework of nonlinear optimization, is [4], "Perturbation of generalized Kuhn-Tucker points in finite-dimensional optimization." Again, the issue is what happens to the optimal solution to a given problem relative to shifts in the parameter values on which the problem depends. When the parameters in question are random variables, this comes down to the study of the statistical distribution of the optimal solution as derived from the distributions of the data elements. Another application of the results is equally important, however. This is to the sensitivity of the optimal policy obtained in the scenario model, discussed earlier, relative to the choice of the probability weights assigned to the branching events in the scenarios. Inasmuch as these weights may in many cases largely be a product of guesswork, it is essential to have a handle on how crucial their values are to an optimal policy determined by computation.

The paper [4] provides a method of testing the effects of alterations in the values. If the effects are large in a given case, this can focus the modeler's attention on a possible trouble spot in the formulation, where perhaps more detail in the scenarios and harder thinking about the assigned probability weights is called for. If the effects are small, on the other hand, the modeler can be reassured that rough guesses are adequate. This can help to justify a particular problem formulation and is a welcome tool therefore in such a difficult modeling area, where one has to cope with uncertainty of many kinds.

Another Ph.D. student, Sien Deng, who will get his degree in the summer of 1993, has worked on this form of approximation in stochastic programming—the sensitivity of stochastic programming problems to the probability values specified with the data. He has written a code to test the sensitivity in two-stage models. This code utilizes Zhu's code [11] as a subroutine.

Yet another student has been done research on the proximal point algorithm in roles related to those in Zhu's work, which seem to be crucial to the methodology of large-scale optimization quite generally. This is Maijian Qian, who finished in August of 1992. Her dissertation [16] provides quasi-Newton schemes for carrying out proximal point iterations to achieve higher rates of convergence. In effect, the geometry of the space is altered from the Euclidean geometry of the canonical norm in order to take advantage of the local geometry generated from a problem's structure around its solution.

Splitting Methods

Still another tack toward the solution of large-scale problems has been taken in [18]. This work, joint between the P.I. and his student George (Hong-gang) Chen, concerns decomposition through "forward-backward splitting." Such splitting, although originally developed for certain kinds of problem decomposition related to boundary value problems involving partial differential equations, has not previously been applied to optimization problems in a Lagrangian format, as is typically advantageous for extended linear-quadratic programming.

We have found that in the case of our problems with dynamic and stochastic structure a surprising and dramatic form of decomposition occurs: it is only necessary repeatedly to solve *small-scale, deterministic* subproblems located in a *single* time period. Less clear yet is what rate of convergence can be obtained numerically in exploiting this idea. Paper [17] is devoted to a series of results on convergence which shed light on the issue, but much more remains to be done, not only theoretically but on the computational front. Chen has been coding the method and will soon have experimental data, which will be included in his dissertation along with the additional theory in the technical reports [19], [20], and [21]. He will experiment with the numerical examples in stochastic programming that Roger Wets and the P.I. have put together for the purpose.

References: Publications Supported by this AFOSR Grant

1. R. Tyrrell Rockafellar and Roger J-B Wets, "Scenarios and policy aggregation in optimization under uncertainty", Math. of Operations Research 16 (1992), 119-147.

2. R. Tyrrell Rockafellar, "Computational schemes for large-scale problems in extended linear-quadratic programming," Math. Programming 48 (1990), 447–474.

3. R. Tyrrell Rockafellar, "Large-scale extended linear-quadratic programming and multistage optimization," in *Advances in Numerical Partial Differential Equations and Optimization* (S. Gomez, J.-P. Hennart, R. Tapia, eds.), SIAM Publications, 1991, 247–261.

4. R. Tyrrell Rockafellar, "Perturbation of generalized Kuhn-Tucker points in finite-dimensional optimization," in *Nonsmooth Optimization and Related Topics*, F. H. Clarke et al. (eds.), Plenum Press, 1989, 393-402.

5. R. Tyrrell Rockafellar, "Generalized second derivatives of convex functions and saddle functions", SIAM J. Control Opt. 28 (1990), 810–822.

6. Alan J. King and R. Tyrrell Rockafellar, "Sensitivity analysis for nonsmooth generalized equations," Math. Programming 55 (1992), 193–212.

7. Alan J. King and R. Tyrrell Rockafellar, "Asymptotic theory for solutions in generalized M-estimation and stochastic programming", accepted for publication in Math. of Operations Research (1993).

8. Stephen E. Wright, *Convergence and Approximation for Primal-dual Methods in Large-scale Optimization*, doctoral dissertation, Department of Mathematics, University of Washington, December, 1990 (99 pages).

9. Stephen E. Wright and R. Tyrrell Rockafellar, "DYNFGM: Dynamic Finite Generation Method," documented FORTRAN code, technical report, Department of Mathematics, University of Washington, March 1991 (revised version).

10. Ciyou Zhu and R. Tyrrell Rockafellar, "Primal-dual gradient projection methods for extended linear-quadratic programming," accepted for publication in SIAM J. Optimization (1993).

11. Ciyou Zhu and R. Tyrrell Rockafellar, "PDCG: Primal-dual Conjugate Gradient Method," documented FORTRAN code, technical report, Dept. of Applied Mathematics, University of Washington, July, 1991.

12. Ciyou Zhu, *Methods for Large-scale Extended Linear-Quadratic Programming*, doctoral dissertation, Dept. of Applied Mathematics, University of Washington, August, 1991 (91 pages).

13. R. Tyrrell Rockafellar and R. J-B Wets, "A dual strategy for the implementation of the aggregation principle in decision making under uncertainty," accepted for publication in *Applied Stochastic Models and Data Analysis*.

14. R. Tyrrell Rockafellar, "Extended linear-quadratic programming," SIAG/OPT Views-and-News, No. 1 (Fall, 1992), 3–6.

15. R. Tyrrell Rockafellar, "Lagrange multipliers and optimality," accepted for publication in SIAM Review (1993).

16. Maijian Qian, *Variable Metric Proximal Point Algorithms*, doctoral dissertation, Dept. of Mathematics, University of Washington, August, 1992 (130 pp.)

17. R. Tyrrell Rockafellar, "Basic issues in Lagrangian optimization," to appear in a proceedings volume published by Physica-Verlag.

18. George H. Chen and R. Tyrrell Rockafellar, "Forward-backward splitting methods in Lagrangian optimization", submitted for publication in SIAM J. Optimization.

19. George H. Chen and R. Tyrrell Rockafellar, "Convergence and structure of forward-backward splitting methods", working paper.

20. George H. Chen and R. Tyrrell Rockafellar, "Application of a splitting algorithm to optimal control and extended linear-quadratic programming", working paper.

21. George H. Chen and R. Tyrrell Rockafellar, " Extended forward-backward splitting methods and convergence", working paper.